arXiv:1608.06175v1 [cs.AI] 17 Aug 2016

# Effectiveness of greedily collecting items in open world games

Andrej Gajduk

***Abstract*— Since *Pokemon Go* sent millions on the quest of collecting virtual monsters, an important question has been on the minds of many people: Is going after the closest item first a time-and-cost-effective way to play? Here, we show that this is in fact a good strategy which performs on average only 7% worse than the best possible solution in terms of the total distance traveled to gather all the items. Even when accounting for errors due to the inability of people to accurately measure distances by eye, the performance only goes down to 16% of the optimal solution.**

## I. INTRODUCTION

The traveling salesman problem (TSP) is a well known NP-hard problem. It asks the following question: *How can a traveling salesman visit all of N cities at least once and return to his starting location in a most efficient manner?* Efficient in this context can refer to distance, time or money. Because of its importance in real world scenarios such as delivery networks and supply chains it has been extensively studied and many optimization algorithms have been proposed in the literature [1], [2], [3]. However, humans have repeatedly been shown to perform better than computer algorithms on specific instances of the problem [4], [5].

Another emerging scenario in which TSP is relevant are open world computer games. These games, sometimes also called free roam, allow players to freely explore a vast virtual world, completing missions and gathering collectibles in their preferred order. Some famous games that fall into this category are *Grand Theft Auto*, *World of Warcraft*, *Assasins Creed* and *Far Cry* featuring modern-urban, fantasy, historic and pacific island settings, respectively. Most recently, *Pokemon GO* used augmented reality to make the whole world a setting for an open world game. All these games feature a range of items, hidden notes, virtual currency stashes or pokemons scattered throughout the virtual world. One of the objectives for the players is to collect as many of these collectibles as possible, especially if they are looking to earn the *Completionist* achievement.

In order to investigate how players of open world games can effectively gather collectibles we focus our attention on a variation of the TSP, where the player must gather all the collectibles starting from a given position, but returning to the starting position is not required. Furthermore, the collectibles are located on a 2D map, and moving between any pair of collectibles is unobstructed as in most games. The optimization parameter is the total distance traveled in order to to collect all the collectibles.

*This work was not supported by any organization
Source code available at https://github.com/gajduk/greedy-tsp

Our goal is to see how a simple greedy algorithm compares to the optimal solution in realistic scenarios. The algorithm is outlined in Section II and our finding related to its performance in Section III. Section IV concludes this manuscript.

## II. ALGORITHM

The algorithm of interest is based on the following gaming strategy: *Always first gather the item which is closest to your current position*. This strategy is what a player would generally use because it is straightforward and does not require a lot of planning.

To find the optimal solution an algorithm based on exhaustive search was used. It considers all possible permutations of the collectibles, and is thus guaranteed to find the optimal solution. Because the number of permutations is proportional to the factorial of the number of collectibles $N!$ this algorithm works only when $N \leq 13$, which subsequently affects our analysis. The **greedy** algorithm works as follows:

**Algorithm 1** Greedy algorithm

```
procedure GREEDY(start_pos, C)
    current_pos ← start_pos
    V ← {}
    while |V| < |C| do
        next_pos ← min_c dist(current_pos, c), ∀c ∈ C \ V
        V ← V ∪ next_pos
        current_pos ← next_pos
    end while
end procedure
```

where $C$ is a set of the positions of all the collectibles, and *dist* is a function which determines the Euclidean distance between two positions in a 2D plane.

Fig. 1. An excerpt from the map from Far Cry 3. Pitchers denote the locations of collectibles

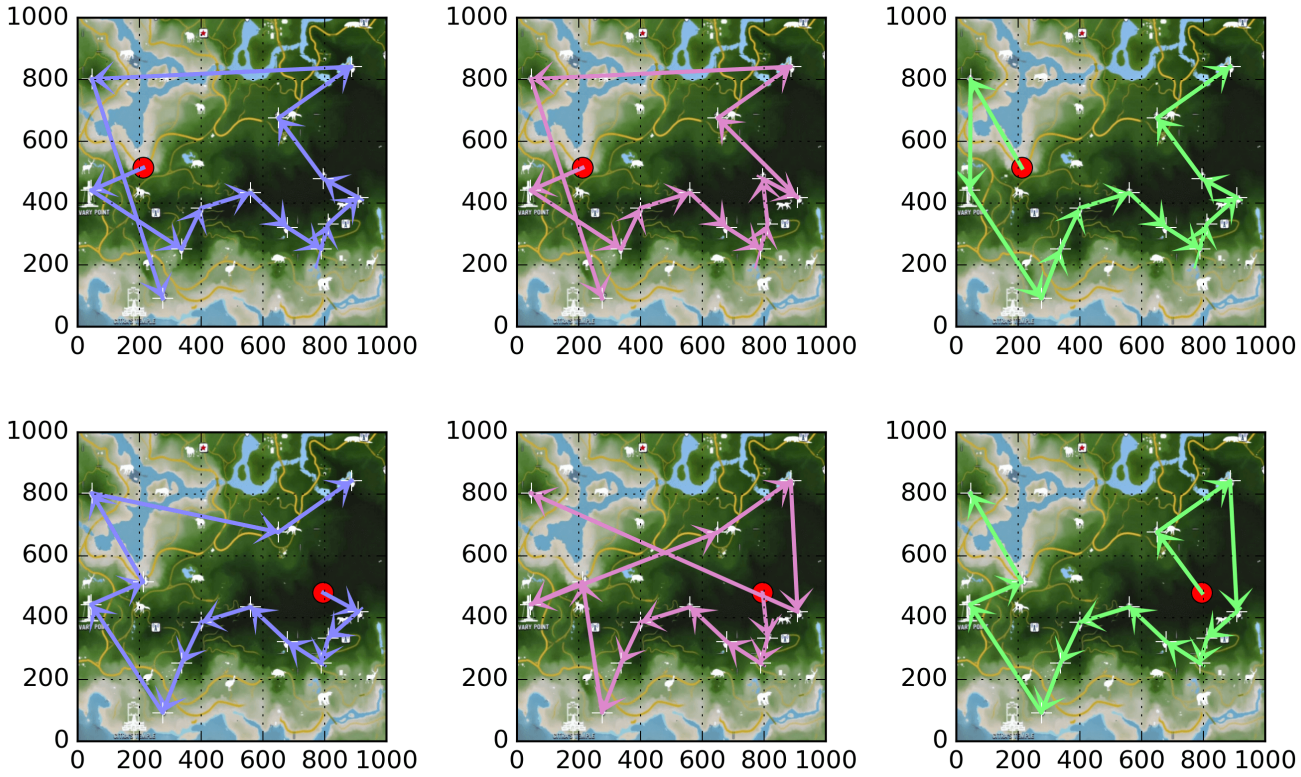


Fig. 2. Rows are different starting positions (red circle). Columns are different algorithms, from left to right: greedy, greedy with error ($\sigma = 0.4$) and optimal. Total distances are for the first row $3601 - 3740 - 2764$; second row $2961 - 3923 - 2883$. The performance of the greedy algorithm depends on the choice of starting position with the difference being much larger for the first row (30.2% longer distance compared to the optimal) as compared to the second row (2.4%). Furthermore the greedy algorithm with error clearly under-performs when compared to the greedy algorithm.

While a player is immersed in the game however, it might not be possible to measure the distance between two collectibles exactly, so the playes has to guestimate it, which may result in errors. To capture this, we propose the **greedy with error** algorithm which includes the following modification to the line 5 in the greedy algorithm

$$next_pos \leftarrow \left[\min_c dist(current_pos, c) + \zeta\right], \forall c \in C \setminus V$$

where $\zeta$ is an error term from a normal distribution $\mathcal{N}(1, \sigma)$ with cutoffs at 0.7 and 1.3 to prevent unreasonable errors.

## III. Results

For illustration, an excerpt from the Far Cry 3 world will be used, where collectibles are marked by pitchers (Fig. 1).

Not surprisingly, the performance of the greedy algorithm depends on the starting location as shown in Figure 2. In the first case, the route chosen by the greedy algorithm is suboptimal and results in a total distance of 3601 (Figa. 2-upper left), opposed to the optimal distance of 2764 (Fig. 2-upper right), which is a difference of 30.2%. In the second case a different starting location is chosen which results in the greedy algorithm finding a solution that is much close to the optimal one ($2961 - 2883$ or a difference of 2.4%), as indicated by the similar routes in the second row of Figure2.

Although, this example is intriguing a detailed analysis is needed to quantify the performance of the greedy algorithm.

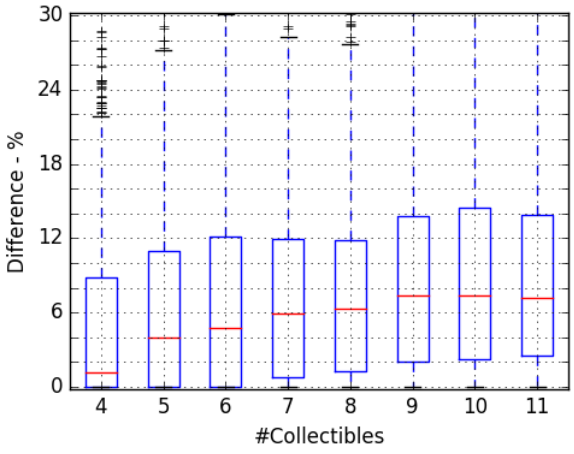


Fig. 3. Average performance of the greedy algorithm compared to the optimal solution, for different number of collectibles.

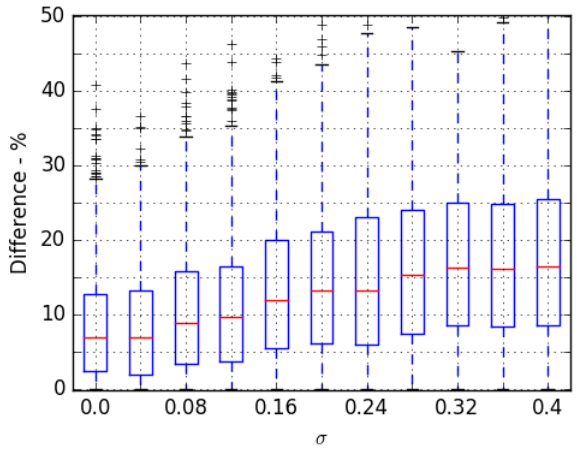


Fig. 4. Average performance of the greedy algorithm with error compared to the optimal solution, for $N = 10$ and different levels of error.

To this end the following setup was used: $N$ collectibles were randomly positioned on a 2D plane (1000x1000 in size), assuming a uniform distribution. The center of the plane with coordinates (500,500) was set as the starting position. Afterwards the optimal distance to get all the collectibles was calculated as well as the total distance using the greedy algorithm. This was repeated 1000 times. Note that, the actual choice of plane size does not influence our results as our only interest is the relative performance of the greedy algorithm as compared to the optimal solution.

We found that the performance of the greedy algorithm deteriorates as the number of collectibles $N$ increases (Fig. 3). However, the overall performance was surprisingly good and for eleven collectibles $N = 11$ it was on average only 7.3% worse than the optimal solution, with upper and lower quartiles at 2.6% and 13.9% respectively.

Next, let us turn our attention to the greedy algorithm with error. This algorithm is not only worse than the optimal, but it is also worse than the greedy algorithm (Fig. 2). Further analysis was performed in order to estimate what effect the level of error in guestimating the distance has on the performance of the greedy algorithm with error. To this end, the number of collectibles was fixed to $N = 10$ and the variance ($\sigma$) in the error term was varied. As expected, higher errors when estimating the distance between two points lead to increasingly poor performance up to 16.9% worse than the optimal (Fig. 4).

## IV. Conclusion

In this manuscript we looked at how a greedy algorithm compares to exhaustive search for a variant of the Traveling Salesman Problem, which is relevant in many open world computer games. The main advantages of the greedy algorithm is that it can be performed by a player as he or she is playing, since no complex computations are required, as opposed to exhaustive search which can only be done by a computer. We found that the performance of the greedy algorithm is comparable to the optimal solution, with the difference depending on the number of collectibles and choice of starting position. For $N = 11$ the greedy algorithm resulted in routes 7.3% longer on average.

Even when human errors in estimating the distances were included the greedy algorithm performed on average less than 16.9% worse than the optimal solution, depending on the level of error. Our analysis was unfortunately limited by the fact that finding the optimal solution is NP hard with complexity $N!$, which prevented us from calculating the performance for $N > 13$. Nevertheless, the results suggest that players who use the greedy strategy of always going after the closest item first, are in fact playing close to optimally in terms of gathering collectibles.

## References

[1] Bellman, Richard. "Dynamic programming treatment of the travelling salesman problem." Journal of the ACM (JACM) 9.1 (1962): 61-63.

[2] Larraaga, Pedro, et al. "Genetic algorithms for the travelling salesman problem: A review of representations and operators." Artificial Intelligence Review 13.2 (1999): 129-170.

[3] Jones, Jeff, and Andrew Adamatzky. "Computation of the travelling salesman problem by a shrinking blob." Natural Computing 13.1 (2014): 1-16.

[4] MacGregor, James N., and Tom Ormerod. "Human performance on the traveling salesman problem." Perception and Psychophysics 58.4 (1996): 527-539.

[5] MacGregor, James N., Chu, Yun, "Human performance on the traveling salesman and related problems: A review", Journal of Problem Solving, 3.2 (2011).